

Dynamic Frequency Scaling and Energy Saving in Quantum Chemistry Applications

Vaibhav Sundriyal
Masha Sosonkina
Fang Liu
Ames Laboratory/DOE
Iowa State University
Ames, IA 50011, USA

{vaibhavs,masha,fangliu}@scl.ameslab.gov

Michael W. Schmidt
Department of Chemistry
Iowa State University
Ames, IA 50011, USA
mike@si.msg.chem.iastate.edu

Abstract—Modern high-performance computing system design is becoming increasingly aware of the energy proportional computing to lower the operational costs and raise reliability. At the same time, high-performance application developers are taking pro-active steps towards less energy consumption without a significant performance loss. One way to accomplish this is to change the processor frequency dynamically during application execution. In this paper, a representative computationally-intensive HPC application GAMESS is considered with the aim to investigate the energy saving potential of its various stages. GAMESS is a quantum chemistry software package used worldwide to perform *ab initio* electronic structure calculations. This paper presents energy consumption characteristics of two Self-Consistent Field method implementations in GAMESS, which radically differ in their computer resource usages. The dynamic frequency scaling optimization is applied to these implementations and serves as verification for the proposed general energy savings model. The developed model provides the minimum of on the compute node energy consumption under a given performance loss tolerance for various processor frequencies.

I. INTRODUCTION

In the past decades, the main focus of High-Performance Computing (HPC) community has been to maximize performance at any cost. The Jaguar Cray XT5 at ORNL (Oak Ridge National Laboratory) has a throughput of about 2.3 petaFLOPS. Consequently, energy and power consumption of these clusters have also increased tremendously. The Green500 list gives the ranking of most power efficient supercomputers on the TOP500 list. The QPACE SFB TR Cluster tops the Green500 list having 773.38 MFLOPS/W with a power consumption of 57.5 kilowatts. Even the Dawning Nebulae, TC3600 blade CB60-G2 cluster, which stands at the fourth position in the Green500, list has a power consumption exceeding 2.5 megawatts.

Dynamic Voltage and Frequency Scaling (DVFS) technique is usually employed to modify the frequency and voltage of the system to save power and energy. The DVFS technique is motivated by the fact that some codes having considerable memory or I/O access, may be executed at a lower frequency without any major impact on their execution time. Owing to the speed difference between processing

element (PE) and memory, PE stalls are generated while waiting for memory response. If frequency is reduced during these stalls, it leads to energy savings without performance degradation.

A common approach for applying DVFS is to perform an optimization of energy for the application [7], [8], [9], to maximize energy savings under a specific user defined performance loss tolerance. In addition, the energy saving of only the PE is often considered, rather than that of the whole compute node (CN), in determining the performance loss tolerance. But a DVFS-based optimization resulting in some energy saving for PE can actually have an overall higher energy consumption for the CN. In this paper, a model is proposed for finding the best case for the overall energy minimization, which may be solved using linear programming techniques. Although the model does not deal with improving the energy efficiency of GAMESS, it depicts how the energy consumption increases with the increase in the performance loss. Hence, to consider only the dynamic power consumption of the PE, as e.g., in [7] and [9], may not be enough for determining overall energy savings.

A. Overview of quantum chemistry package GAMESS

GAMESS [11] is one of the most representative quantum chemistry applications used worldwide to do *ab-initio* electronic structure calculations. GAMESS iteratively approximates the solution to the Schrödinger equation in the form of the Self Consistent Field (SCF) method followed by higher levels of theory, such as Density Functional and Many Body Perturbation. Although GAMESS may be considered as a part of SPECCPU 2006 benchmark suite, studying GAMESS as a stand-alone package yields itself to an investigation of a rich spectrum of quantum chemistry methods and their execution modes. The SCF method is implemented in two forms, namely *direct* and *conventional*, which differ in the handling of the two-electron (2-e) integrals. Specifically, the conventional mode calculates them once at the beginning of the SCF and stores them on disk for subsequent reuse, whereas the direct mode recalculates the 2-e integrals for each iteration. After the SCF, the gradient and Hessian of the

Table I
INPUT SET OF MOLECULES.

Molecule	I/O,(GB)
Silatrane	1.0
Luciferin	3.8
Amg221	5.9
cAMP	7.5
Saxitoxin	9.8
Qinghaousu	11.1
Quinine	13.0
Rotenone	17.1
Ergosterol	22.7

energy may be calculated. In this work, the power and energy characteristics of these two implementations are investigated for a set of molecules.

The rest of paper is organized as follows. Section II describes the power and energy characteristics of the two SCF implementations in GAMESS for a set of input molecules. In Section III, the energy consumption model is derived for varying performance loss tolerances. Section IV experimentally verifies the model. Sections V and VI provide related work and conclusions, respectively.

II. GAMESS ENERGY CHARACTERISTICS

A set of molecules (Table I), is used as inputs to determine the power consumption characteristics of the two SCF implementations. The molecules (column *Molecule*) are listed in the increasing order of their I/O requirements (column *I/O*), as specified in their input files, for the conventional mode. For the experiments, the Ames Lab cluster called “Borges” was used. It consists of four nodes, each having two dual-core 2 GHz Xeon “Woodcrest” CPUs and 8 GB of RAM. The nodes are interconnected with both Gigabit Ethernet and DDR Infiniband. Each processor has a shared 4 MB L2 cache and a 32 KB L1 instruction and data cache per core. Another computing platform (denoted “FScal”) comprises two Dell Optiplex 960 nodes, each of which has an Intel core 2 Duo processor with 2 GB of RAM. This platform was chosen since it allows the CPU frequency scaling for the DVFS optimization. To measure the system power and energy consumption in either platform, a Wattsup power meter [1] was employed.

The power consumption characteristics of GAMESS are explored for a set of input molecules run on Borges. In each node of the four nodes, four processes were executed (denoted as the “4x4” execution configuration), whereas the notation “4x1” stands for executing one process in each node.

A. The 4x4 execution configuration

Fig. 1 depicts the execution time, average power consumption, and energy consumption observed for the various input

molecules in the 4x4 configuration. It can be seen in Fig. 1(a) that, when the I/O requirement is rather low (see Table I), as in the case of Silatrane, the conventional mode performs better than the direct one. However, as the I/O requirement increases, the direct mode begins to outperform greatly the conventional. The situation is exacerbated by a slow I/O rate of around 107 MB/s on Borges. For instance, in another computer with the I/O rate of approximately 226MB/s, the execution time for the conventional mode was almost halved. In spite of the direct mode reducing execution time under the high I/O requirements, conventional mode is more preferable from the computational accuracy point of view and may lead to faster method convergence. Fig. 1(b) depicts the average power consumption for both conventional and direct modes. It can be observed here that the average power consumption of the conventional mode is less than that of direct. This is so because the conventional mode, being I/O intensive, consumes less power as compared with the direct mode, which is more PE and memory intensive. Moreover, as discussed, the PE goes into the idle state quite frequently for the conventional mode, and thus lowering the power consumption. On average, the conventional mode consumes 16% less power as compared with the direct one.

Fig. 1(c) shows that, except for Silatrane, direct mode is more energy efficient than that of conventional in spite of its power consumption being higher. This is mainly due to the fact that the conventional mode suffers from I/O stalls and thus, takes much longer to execute. To summarize, the conventional mode is seemingly more power efficient whereas the direct mode is more energy efficient.

B. The 4x1 execution configuration

Transition to the 4x1 configuration increases the execution time, shown in Fig. 2(a), of the direct mode as compared to the 4x4 direct mode since, being more PE intensive, this mode takes longer to execute on a reduced number of cores. On the other hand, the execution time of the conventional mode decreases in this configuration as compared to the 4x4 configuration. This can be attributed to the fact that when only one process is executing on a single node, the I/O contention is also less and, hence, the execution time is reduced. However, for Ergosterol, the conventional 4x1 execution time is greater than that of the 4x4 configuration since the iteration phase is computationally intensive and thus, takes longer to execute on fewer cores.

The 4x1 average power consumption (Fig. 2(b)) of the conventional mode is almost the same as that for the 4x4 configuration but is lower for Silatrane, which has low I/O requirements. For the direct mode in the 4x1 configuration, the average power is reduced noticeably compared with that in the 4x4 configuration because only one core is active in each node. The energy consumption of the direct mode remains less than that of conventional, as seen in Fig. 2(c). However, the gap between them is reduced because the direct




Figure 1. 4x4 configuration: (a) Execution time, (b) Average power consumption, and (c) Energy consumption normalized with respect to the direct mode for the input molecules shown on the x axes in the ascending order of their I/O requirements.

mode consumes more energy—due to a longer execution—and the conventional consumes less as compared to 4x4 configuration.

C. Power profile of SCF phases

For the conventional mode, the *Calculation* phase is where the 2-e integrals are calculated and stored on the disk. The *Iteration* phase represents the SCF iterations and *Gradient* is where the gradient of the energy is computed. Since the direct mode recomputes integrals for each iteration, it only consists of the *Iteration* and *Gradient* phases. Fig. 3 shows the power profiles of the phases across a node during the entire execution time in the 4x4 configuration.

Fig. 3(a) and Fig. 3(b) depict the power profile of Luciferin. It can be observed that the conventional mode for Luciferin has a very smooth power curve as conventional mode is quite I/O intensive. A molecule executing in the conventional mode suffers from the PE stalls. As a result, PE power consumption is quite low [6] and PE often goes




Figure 2. 4x1 configuration: (a) Execution time, (b) Average power consumption, and (c) Energy consumption normalized with respect to the direct mode for the input molecules shown on the x axes in the ascending order of their I/O requirements.

into the idle state. Therefore, the power consumption of the whole CN is reduced considerably. For the *Gradient* phase, the power consumption increases since it is PE intensive. On the other hand, the power curve for the direct mode has spikes and valleys with power varying between 221 and 255 watts. The PE is idle most of the time during the *Calculation* phase in the conventional mode. Therefore, the average power consumptions of PE and CN conventional modes are quite low comparing with those of the direct mode.

Fig. 3(c) shows the power profile for only the Ergosterol conventional mode. As the I/O requirements of Ergosterol are quite high, the PE spends more time idling as compared with the molecules having lower I/O requirements. Conversely, for Silatrane (Fig. 3(d)), the I/O requirements are low, and the power consumption varies from 225 to 253 watts.

Fig. 4 depicts the power profiles for the molecules exe-

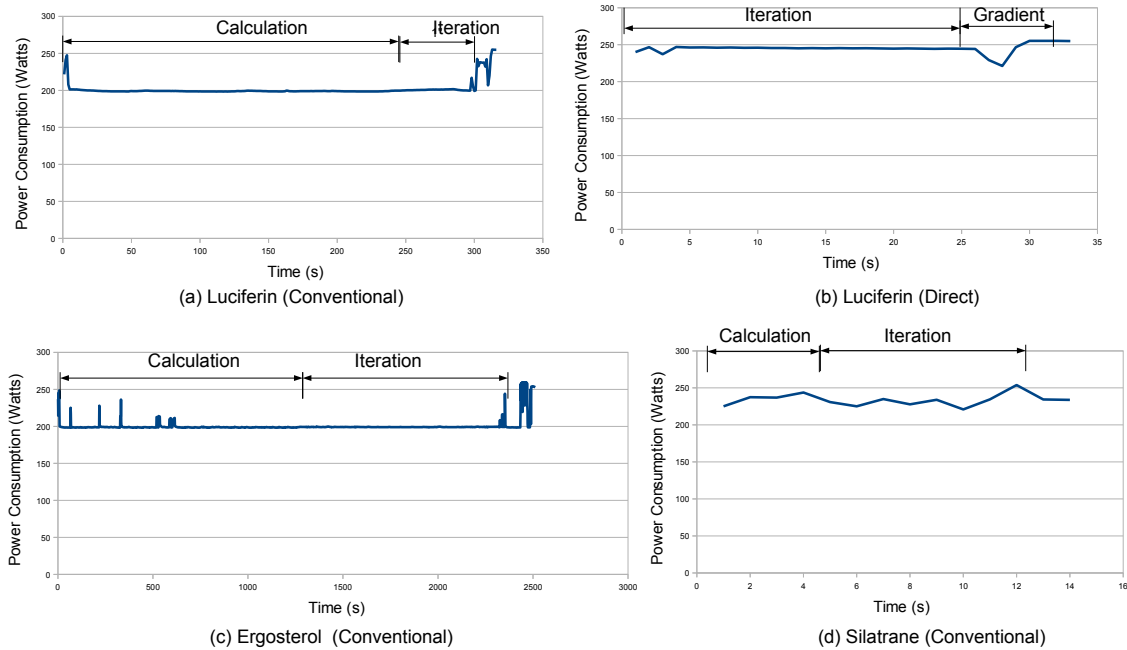


Figure 3. Power profiles for some molecules in the 4x4 configuration.

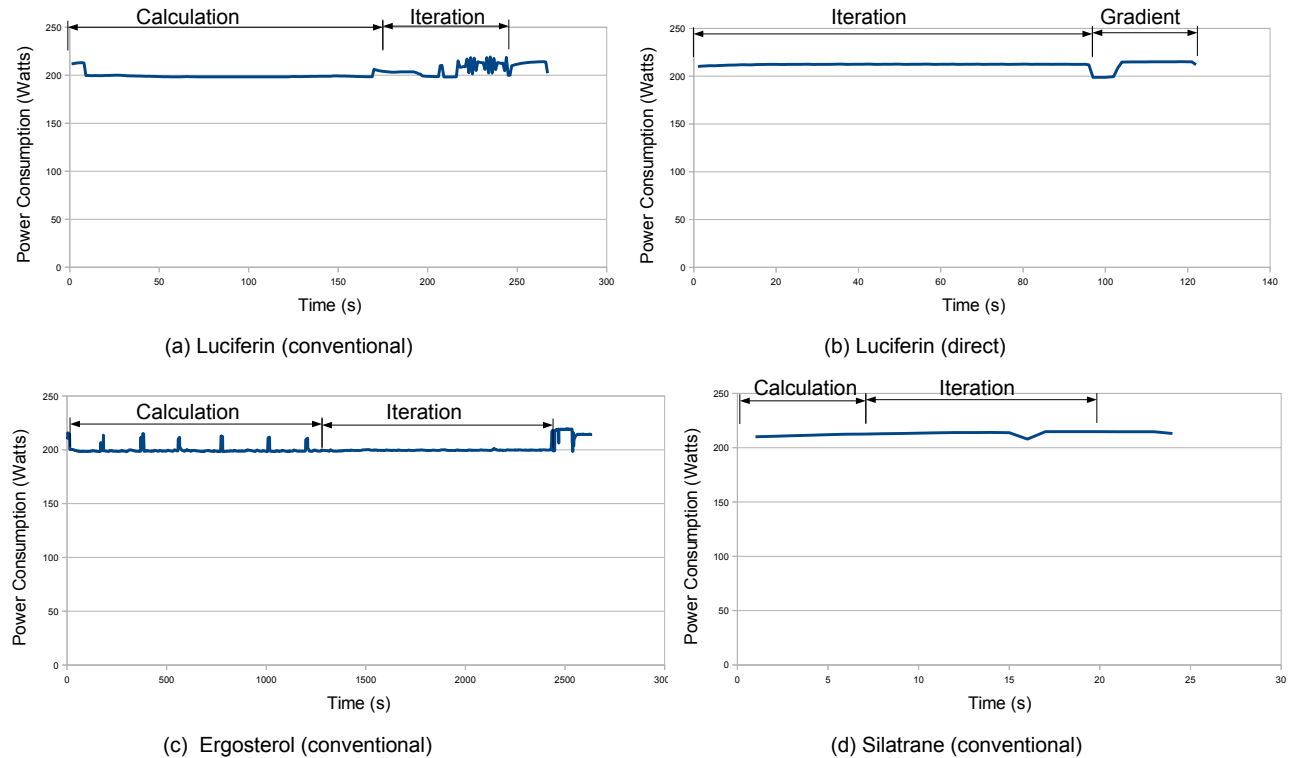


Figure 4. Power profiles for some molecules in the 4x1 configuration.

cuted in the 4x1 configuration. Since the three cores remain idle, the *peak* power consumption is considerably less than that in the 4x4 configuration throughout the experiments. The overall reduction is seen in the direct mode (Fig. 4(b)) which is compute-intensive and suffers when fewer processes are employed. For the conventional mode to which the molecules with high I/O requirements are input (Fig. 4(a) and Fig. 4(c)), there is not much difference in the overall power consumption. On the other hand, for Silatrane with low I/O, the power consumption is also reduced for the conventional mode as seen in Fig. 4(d) comparing with Fig. 3(d).

III. ENERGY CONSUMPTION MODEL

The execution time of a program can be divided into two separate parts, on-chip time t_{on} and off-chip time t_{off} , such that t_{on} and t_{off} are non-overlapping [2]. The time t_{off} consists of stall cycles, such as memory, I/O, branch misprediction, and reservation station stalls, during which the PE is not doing any useful work. In an out-of-order processor, the stall cycles can also overlap with the on-chip execution. DVFS affects only t_{on} of the program execution. For example, if the execution time of a program at the highest frequency f_1 is $t_1 = t_{on} + t_{off}$, then, on a frequency f_i ($f_1 > f_i$), the execution time would be

$$t_i = t_{on}(f_1/f_i) + t_{off} . \quad (1)$$

Typically, during a DVFS-based optimization, a performance loss tolerance is prescribed by the user for a given application, and the energy savings are maximized under this tolerance. Much research focused on applying such an optimization to the PE-only energy savings. However, an optimization resulting in some PE energy savings may actually have a higher overall energy consumption for the whole CN.

Let a DVFS-based optimization increase t_{on} by a factor of k , so that $t' = kt_{on} + t_{off}$. The total energy saving may appear if

$$P_1 t > \bar{P} t' , \quad (2)$$

$$P_1(t_{on} + t_{off}) > \bar{P}(kt_{on} + t_{off}) , \quad (3)$$

where P_1 is the average power consumption of the CN at the highest frequency and \bar{P} is the average power consumption when DVFS is applied. The inequality (3) may be used to determine the feasibility of total energy saving, being rewritten for convenience, as

$$\frac{t_{off}}{t_{on}} > \frac{k\bar{P} - P_1}{P_1 - \bar{P}} . \quad (4)$$

Specifically, Fig. 5 depicts the average power consumption of the input molecules executed in the direct mode on the FScal platform on four available frequencies. Fig. 6 shows the ratio t_{off}/t_{on} (denote it as τ) of the off-chip and on-




Figure 5. Average power consumption of input molecules for direct mode on the four frequencies of FScal platform.




Figure 6. Computation times ratio τ for input molecules normalized with τ for Quinine (FScal platform).

chip computation times for the input molecules normalized with respect to τ of Quinine. It can be observed that the average power consumption varies inversely with τ , which is understandable. As the off-chip accesses increase, the power consumption goes down. Therefore, the average power consumption P_i at a frequency f_i may be written as

$$P_i = a\tau + b , \quad (5)$$

where a and b are some constants, which may be determined using a regression analysis of (5). Their values are shown in Table II along with the correlation coefficient R^2 , which is close to unity. Thus, there is a strong correlation between the ratio τ and the average power consumption at a given frequency.

The PE power consumption consists of two components, static and dynamic. The static power is the power consump-

Table II
REGRESSION COEFFICIENTS FOR DETERMINING THE AVERAGE POWER
CONSUMPTION AT A GIVEN FREQUENCY.

Frequency	b	a	R^2
3.00	74.4	-17.575	0.983
2.67	66	-10.76	0.99
2.33	59	-12.722	0.981
2.00	54	-8.387	0.994

tion of a PE in the idle state. The dynamic power is directly proportional to the product of the operating frequency and the square of the core voltage [10]. In (4), the power consumptions P_1 and \bar{P} can be replaced by the respective PE power consumption values [10] to obtain the feasibility of energy savings for the PE under a DVFS-based optimization.

Since the DVFS affects only the on-chip time, an upper bound of energy savings for a DVFS-based optimization, can be determined. For the maximum energy savings, t_{off} should be executed at the lowest frequency. Therefore, the minimum energy consumption E_{off}^* during the off-chip execution is

$$E_{off}^* = P_n t_{off}, \quad (6)$$

where P_n is the average CN power consumption at the lowest frequency f_n . The performance loss tolerance δ can be defined as

$$\delta = \frac{t' - t}{t} = \frac{(k-1)t_{on}}{t_{on} + t_{off}}, \quad (7)$$

$$k = 1 + \delta(1 + \frac{t_{off}}{t_{on}}). \quad (8)$$

For a user-defined performance loss δ , the execution time t_{on} minimizing the energy consumption on the different frequencies may be determined using linear programming (LP). Let P_1, P_2, \dots, P_n be the average power consumptions of the CN or PE and t_1, t_2, \dots, t_n be the on-chip execution times at frequencies f_1, f_2, \dots, f_n , respectively, where $f_1 > f_2 > \dots > f_n$. By defining f_{ij} as the *scaling factor*, which is the ratio f_i/f_j , the LP problem may be formulated as

$$\min_{t_i} E_{on} = \sum_{i=1}^n P_i f_{1i} t_i, \quad (9)$$

$$\text{such that } \sum_{i=1}^n t_i = t_{on}, \quad (10)$$

$$\sum_{i=1}^n t_i f_{1i} = t_{on}(1 + \delta'), \quad (11)$$

$$t_i \geq 0, \quad (12)$$

where

$$\delta' = \delta(1 + \frac{t_{off}}{t_{on}}). \quad (13)$$

The objective function E_{on} describes the minimum energy consumption of the CN (or PE) during the on-chip work when the DVFS is applied. Note that, during the execution time t_i on the frequency f_i , DVFS changes the average power consumption from P_1 to $P_i f_{1i}$. The formulation (9)–(12) has n variables and two equality constraints and may be solved by using the two-phase Simplex method [12]. In this method, artificial variables are added to the constraints of the original LP and the solution is obtained in two phases in the form of basic and non basic variables. The non basic variables, forming the non basis set, are equal to zero and the basic variables, forming the basis set, provide the optimal solution. For the FScal platform, there are four frequency levels for which the LP (9)–(12) is to be solved here.

Case I: If $0 < 1 + \delta' \leq f_{12}$, the basis set $B = (1, 2)$, the optimal solution is obtained when

$$t_1 = t_{on} - \frac{t_{on}\delta'}{f_{12} - 1} \text{ and } t_2 = \frac{t_{on}\delta'}{f_{12} - 1}.$$

Case II: If $f_{12} < 1 + \delta' < f_{14}$ and $B = (2, 4)$, the optimal solution is obtained when

$$t_2 = \frac{t_{on}(f_{14} - 1 - \delta')}{f_{14} - f_{12}} \text{ and } t_4 = \frac{t_{on}(1 + \delta' - f_{12})}{f_{14} - f_{12}}.$$

Case III: If $f_{14} = 1 + \delta'$, and $B = (4)$, the optimal solution is obtained when

$$t_4 = t_{on}.$$

Once the minimum E_{on}^* is calculated in each of the three cases, the minimum power consumption P^* may be expressed as

$$P^* = \frac{E_{on}^* + E_{off}^*}{t_{on}(1 + \delta') + t_{off}} \quad (14)$$

for each of the three cases considered. This value of P^* can be used in (4) in place of \bar{P} to ascertain the feasibility of energy savings for a DVFS-based optimization since P^* provides a lower bound on \bar{P} (i.e., $\bar{P} \geq P^*$).

IV. MODEL VERIFICATION

Table III shows the PE and CN power consumptions at different frequencies for Saxitoxin direct mode execution. The PE power consumption is determined by using the analytical model proposed in [10], in which authors have accurately determined the static and dynamic power consumption of an Intel Core 2 duo CPU. CN power consumption is calculated by using (3). Using Table III and the E_{on} expression (9),

Table III
CN AND PE POWER FOR SAXITOXIN DIRECT MODE.

Frequency (GHz)	CN Power (W)	Voltage per core (V)	PE Power (W)
3.00	73.4	1.23	52
2.67	65.4	1.20	43
2.33	58.2	1.16	36
2.00	53.4	1.12	28

Table IV
ON-CHIP AND OFF-CHIP TIMES FOR INPUT MOLECULES IN THE FSCAL PLATFORM.

Molecule	t_{on} (s)	t_{off} (s)	τ
Silatrane	19.2	9.2	0.48
Luciferin	74.2	6.6	0.09
Amg221	83.2	9.5	0.11
cAMP	110.0	6.9	0.06
Saxitoxin	119.0	7.0	0.06
Qinghaousu	128.8	15.1	0.12
Quinine	199.6	7.8	0.04
Rotenone	230.4	12.9	0.06
Ergosterol	283.4	20.4	0.07

it may be inferred that, for $i > j$, $P_i f_{1i} > P_j f_{1j}$ in the case of the CN power consumption and $P_i f_{1i} < P_j f_{1j}$ in that for PE. The minimum E_{on}^* is suited for both the CN and PE energy consumption minimizations. However, the CN energy consumption increases with the increase in the performance loss tolerance while the PE energy consumption decreases at the same time.

To determine the times t_{off} and t_{on} for the input molecules, a regression analysis was done on the equation

$$t_i = t_{on} f_{1i} + t_{off}, \quad (15)$$

where t_i is the execution time on frequency f_i , $i = 1, 2, \dots, n$ and t_{on} and t_{off} are constants. The correlation coefficient R^2 ranged from 0.998 to 0.9996 for this regression analysis. Table IV lists t_{off} and t_{on} for the set of input molecules executed in the direct mode on FScal. From the ratio $\tau = t_{off}/t_{on}$, it can be seen that except for Silatrane, direct mode is quite compute intensive. In fact, its execution time scales almost linearly with the change in frequency.

Table V depicts the CN energy consumptions of all the input molecules on the three lower frequencies (columns f_2 , f_3 , f_4) obtained experimentally in the FScal platform. The energy consumption values are normalized with respect to the energy consumption at the highest frequency f_1 . It can be seen that the variation in energy consumption is non-uniform. The theoretical model proposed in Section III may

Table V
CN ENERGY CONSUMPTION IN THE DIRECT MODE ON THREE LOWER FREQUENCIES FOR THE INPUT MOLECULES IN THE FSCAL PLATFORM, NORMALIZED WITH RESPECT TO THE HIGHEST FREQUENCY. FREQUENCIES (IN GHz) ARE $f_1 = 3.0$, $f_2 = 2.67$, $f_3 = 2.33$, AND $f_4 = 2.0$.

Molecule	f_2	f_3	f_4
Silatrane	0.97	0.98	1.01
Luciferin	0.99	1.04	1.09
Amg221	0.99	1.01	1.08
cAMP	0.99	1.04	1.10
Saxitoxin	0.99	1.01	1.03
Qinghaousu	0.98	1.03	1.08
Quinine	1.01	1.04	1.10
Rotenone	1.01	1.04	1.09
Ergosterol	1.01	1.04	1.08

Table VI
DATA SUBSTITUTED INTO THE THEORETICAL MODEL TO DETERMINE THE FEASIBILITY OF THE CN ENERGY CONSUMPTION FOR DIFFERENT FREQUENCIES.

Frequency	rhs	k	\bar{P}
2.00	0.335	1.5	65.4
2.33	0.1	1.288	58.2
2.67	0.01	1.1236	53.4

be used to explain this phenomenon.

Fig. 7 provides an example of the variations in the PE and CN energy consumption at different frequencies for Saxitoxin executed in the direct mode. The values on y axis are normalized with respect to the highest frequency operating point for PE and CN energy, respectively. From the PE energy consumption, it is clear that, as the frequency increases, the energy consumption of PE also increases. This can be verified by putting the values of P_1 and \bar{P} from Table III into inequality (4) for the respective frequencies.

From Table IV, τ ratio of Saxitoxin equals 0.06. To determine the feasibility of the PE energy savings at some frequency, say 2 GHz, the appropriate values may be substituted into inequality (4), such that $P_1 = 52$, $\bar{P} = 28$ (as provided in Table III), and $k = 3/2$. Then, the right-hand side of (4) is equal to -0.714 which is smaller than 0.06, and thus, the PE energy saving is obtained for 2 GHz, as seen in Fig. 7. The same calculation may be performed for other frequencies leading to the conclusion that, for Saxitoxin, the PE energy consumption decreases as the frequency is decreased.

In contrast, the CN energy consumption exhibits non-monotonic behavior for Saxitoxin. It first decreases at 2.67 GHz followed by an increase at lower frequencies. Such a behavior is supported theoretically. In particular, the right-hand side of (4) was calculated for the three lower frequen-




Figure 7. Normalized energy consumption of PE and CN for Saxitoxin in direct mode.

cies, such that their respective \bar{P} values were taken from Table III. The obtained right-hand sides (column *rhs*) as well as the corresponding frequencies (column *Frequency*), average power \bar{P} , and time increase factors k are presented in Table VI while P_1 is fixed at 73.4 for 3 GHz. As shown in Table VI, the condition of CN energy saving is met only at 2.67 GHz, when *rhs* is 0.01. At the other two operating points, 2.33 GHz and 2 GHz, the energy consumption of the CN increases.

Fig. 8 plots the optimum energy consumption E_{on}^* of Saxitoxin (direct mode), as calculated using the model from Section III, for PE and CN versus the performance loss. All the y axis values are normalized with respect to the energy consumption at the highest frequency. The energy consumption for the CN increases with the performance loss increase while the PE is consuming less energy at the same time. Nevertheless, the energy savings are achieved for CN while the performance loss is below 13%, at which point the slope of energy consumption curve increases drastically. These results tally well with the experimental findings for the energy consumption of Saxitoxin shown in Fig. 7. For example, the CN energy consumptions at 2.33 GHz and 2 GHz are no less than that at the highest frequency while energy savings are achieved above 2.67 GHz. The values of the performance loss tolerance δ are, respectively, 27.15%, 47.2%, and 11.68%, marked as vertical straight lines in Fig. 8.

V. RELATED WORK

The DVFS technique has been studied extensively in the past and was applied to the performance optimization with respect to energy or power. A regression-based β -adaptation algorithm is introduced in [7] for the optimization. In [8], the authors characterize stall cycles through the hardware




Figure 8. Variation of energy consumption with performance loss for Saxitoxin in direct mode.

performance counters and save energy during the stalls. The work in [4] studies different distributed DVFS scheduling strategies and their impact on application performance. A performance model and a history-based workload prediction strategy are proposed for energy-performance optimization in [5]. Researchers have also studied different types of distributed applications to find opportunities of energy savings. The power consumption characteristics of NAS parallel benchmark in terms of various CPU components are studied in [6]. In [3], the authors present a framework for automatic profiling of power consumption for parallel scientific applications on high-performance distributed systems. The work in [10] proposes to model the DVFS overhead in the form of execution time penalty and energy consumption.

VI. CONCLUSIONS AND FUTURE WORK

This work studies the power and energy consumption characteristics of GAMESS when performing SCF calculations in two different ways, direct and conventional. By considering them in stages and using different numbers of processing elements (PEs), it has been observed that the direct mode is more energy efficient although its performance suffers when fewer cores are used and has less potential for the DVFS optimization due to the less time spent off-chip. GAMESS is one of the prominent quantum chemistry applications and enjoys a large user community. Hence, the findings of this paper are beneficial to an important scientific domain and a wide class of HPC applications.

A general theoretical model for evaluating a DVFS-based optimization is proposed and verified experimentally using GAMESS. The model demonstrates that the energy consumption for the on-chip time, increases with the performance loss increase. Therefore, care must be taken when choosing a performance loss tolerated for energy savings. The ratio of the off-chip and on-chip execution is critical

in determining the performance loss. The results emphasize that applying DVFS may actually lead to a higher total energy consumption as compared with always keeping the highest frequency. More research is needed, however, to fine tune the model. The memory and I/O power changes with respect to the access patterns need to be studied thoroughly. The goal is to make the application energy-efficient with little or no degradation in its performance.

ACKNOWLEDGMENT

This work was supported in part by Iowa State University under the contract DE-AC02-07CH11358 with the U.S. Department of Energy, by the Director, Office of Science, Division of Mathematical, Information, and Computational Sciences of the U.S. Department of Energy under contract number DE-AC02-05CH11231, and by the National Science Foundation grants NSF/OCI – 0749156, 0941434, 1047772.

REFERENCES

- [1] <https://www.wattsupmeters.com>.
- [2] K. Choi, R. Soma, and M. Pedram. Fine-grained dynamic voltage and frequency scaling for precise energy and performance tradeoff based on the ratio of off-chip access to on-chip computation times. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, 24(1):18 – 28, 2005.
- [3] X. Feng, R. Ge, and K.W. Cameron. Power and energy profiling of scientific applications on distributed systems. In *Parallel and Distributed Processing Symposium, 2005. Proceedings. 19th IEEE International*, page 34, 2005.
- [4] R. Ge, X. Feng, and K.W. Cameron. Performance-constrained distributed DVS scheduling for scientific applications on power-aware clusters. In *Supercomputing, 2005. Proceedings of the ACM/IEEE SC 2005 Conference*, page 34, nov 2005.
- [5] R. Ge, X. Feng, W. Feng, and K.W. Cameron. CPU MISER: A performance-directed, run-time system for power-aware clusters. In *Parallel Processing, 2007. ICPP 2007. International Conference on*, page 18, 2007.
- [6] R. Ge, X. Feng, S. Song, H. C. Chang, D. Li, and K.W. Cameron. PowerPack: Energy profiling and analysis of high-performance systems and applications. *Parallel and Distributed Systems, IEEE Transactions on*, 21(5):658 –671, May 2010.
- [7] C. H. Hsu and W. Feng. A power-aware run-time system for high-performance computing. In *Supercomputing, 2005. Proceedings of the ACM/IEEE SC 2005 Conference*, page 1, nov 2005.
- [8] S. Huang and W. Feng. Energy-efficient cluster computing via accurate workload characterization. In *Cluster Computing and the Grid, 2009. CCGRID '09. 9th IEEE/ACM International Symposium on*, pages 68 –75, May 2009.
- [9] S. J. Lee, H. K. Lee, and P. C Yew. Runtime performance projection model for dynamic power management. In Lynn Choi, Yunheung Paek, and Sangyeun Cho, editors, *Advances in Computer Systems Architecture*, volume 4697 of *Lecture Notes in Computer Science*, pages 186–197. Springer Berlin / Heidelberg, 2007. 10.1007/978-3-540-74309-5_19.
- [10] J. Park, D. Shin, N. Chang, and M. Pedram. Accurate modeling and calculation of delay and energy overheads of dynamic voltage scaling in modern high-performance microprocessors. In *Low-Power Electronics and Design (ISLPED), 2010 ACM/IEEE International Symposium on*, pages 419 – 424, 2010.
- [11] M. W. Schmidt, K. K. Baldrige, J. A. Boatz, S. T. Elbert, M. S. Gordon, J. H. Jensen, S. Koseki, N. Matsunaga, K. A. Nguyen, S. Su, T. L. Windus, M. Dupuis, and J. A. Montgomery, Jr. General atomic and molecular electronic structure system. *J. Comput. Chem.*, 14:1347–1363, November 1993.
- [12] R. J. Vanderbei. *Linear Programming: Foundations and Extensions*. Kluwer Academic Publishers, 2001.